



WLCG Experiments Test Framework Update

Marian Babik, CERN
WLCG operations coordination, Nov 2020



Outline

- Follow up on ETF update [presented](#) at WLCG ops coordination last year
- Today I will focus on the recent major ETF updates
 - Already in production for CMS, ATLAS and soon for others
- New job submission and WN test frameworks
- Quick walkthrough of the new features and functionality
- Challenges and Plans



Motivation

Current WLCG Service Availability Monitoring (SAM) structure

- **WLCG Experiments Test Framework (ETF)**

- WLCG testing middleware (running so called SAM tests)
- Active testing of the sites services and reporting back to SAM3/MONIT
- Common to all experiments
- Main source for WLCG Availability/Reliability Reports (different from EGI/ARGO)

- **MONIT/SiteMon**

- Aggregation (via custom algos), visualisation and reporting
- Support for multiple sources of metrics (e.g. ALICE storage tests, ATLAS ASAP)

- **A generic test framework remains fundamental for WLCG monitoring**

- Keeping track of sites availability/reliability
- Running deployment campaigns (IPv6, HTTP, etc.)
- Provides means of isolation when debugging site/experiments issues
 - Middleware bugs, site setup/configuration, latency sensitive issues/timeouts, etc.
- Contributing to the operational toolchain of the experiments

Overview

Generic test middleware based on open source

- Checkmk, Nagios core and Messaging (ActiveMQ)
- Containerised, uses gitlab CI/Auto DevOps

Focuses on functional testing (atomic)

- Direct job submissions, worker node env. testing
- Core storage operations
- Remote API testing and/or network testing (ping/icmp)

~ 150 sites, 1200 hosts monitored

~ 10 metrics/host

~ 1M metrics/day

ALICE, ATLAS, CMS, LHCb, **DUNE**

High-level functional testing

Site notifications

Checkmk dashboard to show results

Raw 1.5.0p11

Host search pic.es

21 rows /DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=mbabik/CN=555091/CN=Marian Babik (admin) 15:36

Tactical Overview

Hosts	Problems	Unhandled	State
1038	0	0	0

Services

Services	Problems	Unhandled	State
11715	1362	1352	900

Events

Events	Problems	Unhandled	State
0	0	0	0

Quicksearch

pic.es

Views

- Overview
- Host & Services Problems
- Main Overview
- Network Topology
- Hosts
 - All hosts
 - All hosts (Mini)
 - All hosts (tiled)
 - Favorite hosts
 - Host search
- Host Groups
 - Host Groups
 - Host Groups (Grid)

ATLAS

state	Host	Icons	OK	Wa	Un	Cr	Pd
UP	ce13.pic.es		5	0	0	0	0
UP	lfæce04.pic.es		5	0	0	0	0
UP	webdav-at1.pic.es		29	0	0	0	0
UP	xrootd-at2-door.pic.es		16	0	0	0	0

CMS

state	Host	Icons	OK	Wa	Un	Cr	Pd
UP	ce12.pic.es		12	0	0	1	0
UP	srmcms.pic.es		9	0	0	0	0

LHCb

state	Host	Icons	OK	Wa	Un	Cr	Pd
UP	ce13.pic.es		9	0	0	0	0
UP	srmhcb.pic.es		1	0	0	0	0

Host Search Results (pic.es)

state	Host	Icons	OK	Wa	Un	Cr	Pd
UP	ce14.pic.es		5	0	0	0	0
UP	srmatlas.pic.es		16	0	0	0	0
UP	webdav-at2.pic.es		29	0	0	0	0
UP	gridftp.pic.es		16	0	0	0	0
UP	srmifæe.pic.es		16	0	0	0	0
UP	xrootd-at1-door.pic.es		16	0	0	0	0
UP	ce13.pic.es		13	0	0	0	0
UP	xrootd01-cmst1.pic.es		4	0	0	0	0
UP	ce14.pic.es		13	0	0	0	0
UP	xrootd02-cmst1.pic.es		4	0	0	0	0
UP	ce14.pic.es		9	0	0	0	0
UP	srmhcb-tape.pic.es		1	0	0	0	0
UP	webdav-lhcb1.pic.es		13	0	0	0	0


refresh: 60 secs

Plugins/Tests

Plugins	Users/Experiments	Maintained by
<u>Job Submission</u>		
CREAM, ARC, HTCONDOR-CE <i>JESS</i>	LHCb, ALICE, ATLAS, CMS, DUNE	ETF
<u>Worker Nodes</u>		
ATLAS (3), CMS (11), LHCb (7), DUNE(1)	ATLAS, CMS, LHCb, DUNE	ATLAS, CMS, LHCb, DUNE
<u>Storage</u>		
GFAL2 (SRM, gsiftp, XRoot, HTTP)	ATLAS	ATLAS
GFAL2 (SRM)	CMS	CMS
XRoot	CMS	CMS
HTTPs/WebDAV	HTTP TF	HTTP TF
<u>Network</u>		
perfSONAR (14)	WLCG Network Throughput WG	OSG, WLCG

Quick recap - ETF submissions and WN tests

- Job submission
 - **ETF submits one job at a time per CE and follows it up**
 - Each experiment has production and QA instances (so queue can have multiple jobs)
 - **Experiments now in control on most of the aspects related to submission**
 - Which CEs are tested, frequency of job submission, credentials, timeouts, resource constraints, etc.
 - Submission methods specific to each experiment
 - Direct (LHCb, ALICE) or via local HTCondor pool (ATLAS, CMS, DUNE)
 - HTCondor pool is run locally as part of ETF - each experiment has its own
- WN tests
 - **Part of the job payload is the worker node test framework together with all WN tests**
 - Test framework runs all WN tests in parallel (configurable, usually in 2 processes)
 - WN tests are scripts compliant with the nagios plugins standard
 - WN test results are retrieved as part of the output once the job finishes
 - High-level timeout for all tests to execute is 10 minutes, per test timeout is a bit less



New Job Submission Framework (JESS)

- **New job submission [library](#) and new job submission plugin**
 - The aim was to simplify the design and develop a reusable library (not tied to Nagios)
- **Pluggable** - easy to extend to support different submission systems
 - Direct submission to ARC, CREAM and HT-Condor-CE
 - Submissions via local HT-Condor
 - to ARC/CREAM/HTCondorCE and potentially other backends
 - **Submissions via remote HT-Condor pool**
 - ETF can also host a **local HT-Condor pool** to which remote startds can connect
- **Job tracking and monitoring**
 - Currently tracking a single job per CE; can also be extended to track multiple jobs/CE
 - Manual re-scheduling of job submissions via web interface improved
 - Full log(s) of the running job (common structure across different backends)
- **Support for configuration/env on the worker nodes**
 - This can be configured by the experiments in the ETF core plugin (per host/service/site)
- Drops output limit on the submission details and worker node results

Job submission and WN tests

arc-ce05.gridpp.rl.ac.uk

State	Service	Icons	Status detail	Age	Checked	Perf-O-Meter
OK	org.cms.WN-analysis-/cms/Role=lcgadmin		cmssgm-846658.0-lcg2229.gridpp.rl.ac.uk: OK	201 m	36 m	
OK	org.cms.WN-basic-/cms/Role=lcgadmin		cmssgm-846658.0-lcg2229.gridpp.rl.ac.uk: OK	2020-10-22 08:46:25	36 m	
OK	org.cms.WN-cvmfs-/cms/Role=lcgadmin		cmssgm-846658.0-lcg2229.gridpp.rl.ac.uk: OK, cvmfs vers 2.6.0 (probe 1.3-pre2)	2020-10-20 10:01:10	36 m	
OK	org.cms.WN-env-/cms/Role=lcgadmin		cmssgm-846658.0-lcg2229.gridpp.rl.ac.uk: OK	2020-10-20 10:01:10	36 m	
OK	org.cms.WN-frontier-/cms/Role=lcgadmin		cmssgm-846658.0-lcg2229.gridpp.rl.ac.uk: OK	2020-10-20 10:01:10	36 m	
OK	org.cms.WN-isolation-/cms/Role=lcgadmin		cmssgm-846658.0-lcg2229.gridpp.rl.ac.uk: OK	2020-10-20 10:01:10	36 m	
OK	org.cms.WN-mc-/cms/Role=lcgadmin		cmssgm-846658.0-lcg2229.gridpp.rl.ac.uk: OK	2020-10-29 21:46:32	36 m	
OK	org.cms.WN-squid-/cms/Role=lcgadmin		cmssgm-846658.0-lcg2229.gridpp.rl.ac.uk: OK	2020-10-22 08:46:25	36 m	
OK	org.cms.WN-xrootd-access-/cms/Role=lcgadmin		cmssgm-846658.0-lcg2229.gridpp.rl.ac.uk: OK	201 m	36 m	
OK	org.cms.WN-xrootd-fallback-/cms/Role=lcgadmin		cmssgm-846658.0-lcg2229.gridpp.rl.ac.uk: OK	231 m	36 m	
OK	org.sam.CONDOR-JobState-/cms/Role=lcgadmin		OK - Existing job (7054939) was found in status RUNNING	2020-10-20 09:46:08	6 m	
OK	org.sam.CONDOR-JobSubmit-/cms/Role=lcgadmin		Reschedule check BI Aggregations containing this Service Parameters for this service	2020-10-20 10:01:10	36 m	

JobSubmit metric will show the result of the job submission (once job has reached its final state) - this will include all details as in JobState but will also have additional details on job output, timeouts, etc.

WN metrics report results of the WN tests - they're all set by the JobState once the job finishes

JobState metric is the one that submits the job and follows it up - it reports current status of the running job including job log, last status details, JDL, etc. This metric will report all the other metrics for this hosts. JobState is not sent to MONIT.

This is the only metric that can be re-scheduled in order to fetch latest status or resubmit a new job (this will only work if previous job has already finished)

Job details (HTCondor)

arc-ce04.gridpp.rl.ac.uk

org.sam.CONDOR-JobState-/cms/Role=lcgadmin



OK

OK - Existing job (7054569) was found in status RUNNING

=== Job JDL:

```
JDL(((universe, 'grid'), ('executable', 'etf_run.sh'), ('transfer_executable', 'true'), ('output', '/var/lib/gridprobes/cms.Role.lcgadmin/scondor/arc-ce04.gridpp.rl.ac.uk/out/gridjob.err'), ('log', '/var/lib/gridprobes/cms.Role.lcgadmin/scondor/arc-ce04.gridpp.rl.ac.uk/out/gridjob.log'), ('log_', ('use_x509userproxy', 'true'), ('grid_resource', 'nordugrid arc-ce04.gridpp.rl.ac.uk'), ('arguments', '-v cms -c arc-ce04.gridpp.rl.ac.uk -p 2 -t ce04.gridpp.rl.ac.uk/wlogs.tgz', '/var/lib/gridprobes/cms.Role.lcgadmin/scondor/arc-ce04.gridpp.rl.ac.uk/gridjob.tgz')), ('transfer_output_files', 'wlogs.tgz=/var/lib/gridprobes/cms.Role.lcgadmin/scondor/arc-ce04.gridpp.rl.ac.uk/out/wlogs.tgz'), ('should_transfer_files', 'YES'), ('no
```

=== Job submission:

```
condor_submit /var/lib/gridprobes/cms.Role.lcgadmin/scondor/arc-ce04.gridpp.rl.ac.uk/gridjob.jdl
```

Submitting job(s).

1 job(s) submitted to cluster 7054569.

=== Job log:

000 (7054569.0) 2020-11-03T10:01:29 SubmitEvent

SubmitHost:<188.184.104.127:9618?addr=188.184.104.127-9618&noUDP&sock=12_cf39_4>

027 (7054569.0) 2020-11-03T10:01:35 GridSubmitEvent

GridResource:nordugrid arc-ce04.gridpp.rl.ac.uk

GridJobId:nordugrid arc-ce04.gridpp.rl.ac.uk ikGLDmLJ8txnE6QDjqm6UqoABFKDmABFKDmMhMXWdMABFKDm6mDqQm

001 (7054569.0) 2020-11-03T10:05:35 ExecuteEvent

ExecuteHost:nordugrid arc-ce04.gridpp.rl.ac.uk

=== Last job status:

Arguments = "-v cms -c arc-ce04.gridpp.rl.ac.uk -p 2 -t 600 -T 550 -d"

BufferBlockSize = 32768

BufferSize = 524288

ClusterId = 7054569

Cmd = "/var/lib/gridprobes/cms.Role.lcgadmin/scondor/arc-ce04.gridpp.rl.ac.uk/etf_run.sh"

CommittedSlotTime = 0

CommittedSuspensionTime = 0

CommittedTime = 0

CompletionDate = 0

CondorPlatform = "\$CondorPlatform: X86_64-CentOS_7.8 \$"

CondorVersion = "\$CondorVersion: 8.8.10 Aug 12 2020 PackageID: 8.8.10-1.3 \$"

CoreSize = 0

CumulativeRemoteSysCpu = 0.0

JobState metric output is common to all backends and has the following structure (structure changes a bit depending on state of the job):

- Job JDL
- Job Submission - shows submission command and output
- Job log - shows last job log
- Last job status - shows latest job status
- Once completed will also show
 - ETF job log - timeout limits configured and how they were tracked by ETF - times shown are approx - snapshots when JobState metric was executed
 - WN test results - status of parsing from the job's output tarball and ETF job log

CONDOR backend details:

- Submission is done via condor_submit to local or remote HTCondor pool (this example uses local)
- Job log is taken from HTCondor pool log file
- Last job status is taken from condor_q -l command
- JDL will contain resource constraints as defined by the experiment (queues, memory, walltime, etc.)

Job details (ARC and CREAM-CE)

arc-ce01.gridpp.rl.ac.uk
org.sam.ARC-JobState-alice



OK

OK - Job successfully completed (status:Finished, id:gsiftp://arc-ce01.gridpp.rl.ac.uk:2811/jobs/1kvNDmWP8tnCIXDjqlBL5XqABFKDmABFKDm4hrQDmABFKDmmRurtm

=== ETF job log:

Timeout limits configured were:

global -> 1410 minutes

Queueing -> 1380 minutes

Current time: 2020-11-03 10:23:07

Job started: 2020-11-03 10:08:00

Job finished: 2020-11-03 10:22:57

Job tracking times (entered):

Finished -> 2020-11-03 10:22:57

=== Job JDL:

JDL(["(executable','etf_dummy.sh'),('join','yes'),('stdout','arc.out'),('queue','grid3000M'),('runtimeenvironment','ENV/PROXY'),('cputime','1800'),('walltime','1800)"])

=== Job submission:

arcsub -debug INFO --cluster arc-ce01.gridpp.rl.ac.uk --timeout 120 --joblist /var/lib/gridprobes/alice/arc/arc-ce01.gridpp.rl.ac.uk/jobs.dat /var/lib/gridprobes/alice/arc/arc-

INFO: Configuration (/etc/arc/client.conf) loaded

INFO: Configuration (/omd/sites/etf/arc/client.conf) loaded

INFO: Using proxy file: /opt/omd/sites/etf/etc/nagios/globus/userproxy.pem-alice

INFO: Using CA certificate directory: /etc/grid-security/certificates

INFO: Broker Random loaded

INFO: Computing endpoint https://arc-ce01.gridpp.rl.ac.uk:443/arex (type org.nordugrid.arcrest) added to the list for submission brokering

INFO: Computing endpoint https://arc-ce01.gridpp.rl.ac.uk:443/arex (type org.ogf.glue.emies.activitycreation) added to the list for submission brokering

INFO: Computing endpoint gsiftp://arc-ce01.gridpp.rl.ac.uk:2811/jobs (type org.nordugrid.gridftpjob) added to the list for submission brokering

INFO: Transfer from file:/var/lib/gridprobes/alice/arc/arc-ce01.gridpp.rl.ac.uk/etf_dummy.sh to gsiftp://arc-ce01.gridpp.rl.ac.uk:2811/jobs/1kvNDmWP8tnCIXDjqlBL5XqA

INFO: Real transfer from file:/var/lib/gridprobes/alice/arc/arc-ce01.gridpp.rl.ac.uk/etf_dummy.sh to gsiftp://arc-ce01.gridpp.rl.ac.uk:2811/jobs/1kvNDmWP8tnCIXDjqlBL5

INFO: Using buffered transfer method

INFO: write_thread: get and pass buffers

INFO: [external] Using proxy file: /opt/omd/sites/etf/etc/nagios/globus/userproxy.pem-alice

INFO: [external] Using CA certificate directory: /etc/grid-security/certificates

Job submitted with jobid: gsiftp://arc-ce01.gridpp.rl.ac.uk:2811/jobs/1kvNDmWP8tnCIXDjqlBL5XqABFKDmABFKDm4hrQDmABFKDmmRurtm

=== Job log:

No warning: Job not found in job list: gsiftp://arc-ce01.gridpp.rl.ac.uk:2811/jobs/1kvNDmWP8tnCIXDjqlBL5XqABFKDmABFKDm4hrQDmABFKDmmRurtm

Warning found, try later

=== Last job status:

Job: gsiftp://arc-ce01.gridpp.rl.ac.uk:2811/jobs/1kvNDmWP8tnCIXDjqlBL5XqABFKDmABFKDm4hrQDmABFKDmmRurtm

State: Finished

Specific state: FINISHED

Exit Code: 0

Owner: /DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=maarten/CN=410032/CN=Maarten Litmaath

Other Messages: SubmittedVia=org.nordugrid.gridftpjob

Queue: grid3000M

Requested Slots: 1

Stdin: /dev/null

Stdout: arc.out

Stderr: arc.out

Submitted: 2020-11-03 10:07:58

End Time: 2020-11-03 10:09:28

Submitted from: 137.138.62.91:8860

Requested CPU Time: 30 hours

ce1.ts.infn.it

org.sam.CREAMCE-JobState-alice



OK

OK - Job successfully completed (status:DONE-OK, id:https://ce1.ts.infn.it:8443/CREAM895682966)

=== ETF job log:

Timeout limits configured were:

IDLE -> 1380 minutes

global -> 1410 minutes

Current time: 2020-11-03 10:26:46

Job started: 2020-11-03 10:11:42

Job finished: 2020-11-03 10:26:40

Job tracking times (entered):

DONE-OK -> 2020-11-03 10:26:40

=== Job JDL:

JDL(["(OutputSandBox','cream.out','wnlogs.tgz'),('Executable','etf_dummy.sh'),('StdError','cream.out'),('StdOutput','cream.out'),('/var/lib/gridprobes/alice/cream/ce1.ts.infn.it/etf_dummy.sh','/var/lib/gridprobes/alice/cream/ce1.ts.infn.it/gridjob.tgz'),('JobType','Non

=== Job submission:

gLite-ce-job-submit --autm-delegation --debug -r ce1.ts.infn.it:8443/cream-lsf-alice /var/lib/gridprobes/alice/cream/ce1.ts.infn.it/gridjob.tgz

2020-11-03 10:11:40,409 DEBUG - Using certificate proxy file /opt/omd/sites/etf/etc/nagios/globus/userproxy.pem-alice

2020-11-03 10:11:40,425 DEBUG - VO from certificate=[alice]

2020-11-03 10:11:40,425 WARN - No configuration file suitable for loading. Using built-in configuration

2020-11-03 10:11:40,425 DEBUG - Logfile is [/tmp/gLite-cream_cli_logs/gLite-ce-job-submit_CREAM_1002_20201103-101140.log]

2020-11-03 10:11:40,426 DEBUG - Processing file [/var/lib/gridprobes/alice/cream/ce1.ts.infn.it/etf_dummy.sh]...

2020-11-03 10:11:40,427 DEBUG - Processing file [/var/lib/gridprobes/alice/cream/ce1.ts.infn.it/gridjob.tgz]...

2020-11-03 10:11:40,427 DEBUG - Inserting mangled InputSandBox in JDL: [{"/var/lib/gridprobes/alice/cream/ce1.ts.infn.it/etf_dummy

2020-11-03 10:11:40,431 INFO - certUtil:generateUniqueID() - Generated DelegationID: [91d244417b98a678e213cd60d0d3ab034d

2020-11-03 10:11:41,496 DEBUG - Registering to [https://ce1.ts.infn.it:8443/cream/services/CREAM2] JDL=[InputSandBox = {

"/var/lib/gridprobes/alice/cream/ce1.ts.infn.it/etf_dummy.sh","/var/lib/gridprobes/alice/cream/ce1.ts.infn.it/gridjob.tgz"}]; BatchSystem = {

OutputSandBoxBaseDestUri = "gsiftp://localhost"; Executable = "etf_dummy.sh"; QueueName = "alice"; OutputSandBox = { "cream.out

2020-11-03 10:11:41,796 DEBUG - JobID=[https://ce1.ts.infn.it:8443/CREAM895682966]

2020-11-03 10:11:41,796 DEBUG - UploadURL=

[gsiftp://ce1.ts.infn.it/var/cream_sandbox/alice/CN_Maarten_Litmaath_CN_410032_CN_maarten_OU_Users_OU_Organic_Units_DC

2020-11-03 10:11:41,798 INFO - Sending file

[gsiftp://ce1.ts.infn.it/var/cream_sandbox/alice/CN_Maarten_Litmaath_CN_410032_CN_maarten_OU_Users_OU_Organic_Units_DC

2020-11-03 10:11:42,378 INFO - Sending file

[gsiftp://ce1.ts.infn.it/var/cream_sandbox/alice/CN_Maarten_Litmaath_CN_410032_CN_maarten_OU_Users_OU_Organic_Units_DC

2020-11-03 10:11:42,413 DEBUG - Will invoke JobStart for JobID [CREAM895682966]

https://ce1.ts.infn.it:8443/CREAM895682966

=== Last job status:

***** JobID=[https://ce1.ts.infn.it:8443/CREAM895682966]

Current Status = [DONE-OK]

Working Dir = [[reserved]]

ExitCode = [0]

Grid JobID = [N/A]

LRMS Abs JobID = [[reserved]]

LRMS JobID = [[reserved]]

Deleg Proxy ID = [91d244417b98a678e213cd60d0d3ab034d92f88]

DelegProxyInfo = [{"isRFC=true"; valid from="11/3/20 10:06 AM (GMT)"; valid to="11/4/20 09:09 AM (GMT)"; holder DN="CN=Maarten

issuer="CN=1576233846"; valid from="1324024663,CN=1982085939,CN=2015585009,CN=207747510,CN=Maarten Litmaath,CN=410032,CN

OU=computers,DC=cern,DC=ch"; VOMS attributes=(/alice/Role=NULL/Capability=NULL, /alice/alarm/Role=NULL/Capability=NULL,

Worker Node = [farm052.ts.infn.it]

Local User = [alice045]

CREAM ISB URI = [gsiftp://ce1.ts.infn.it/var/cream_sandbox/alice/CN_Maarten_Litmaath_CN_410032_CN_maarten_OU_Users_OU_


Plugin output

Historically plugin output has been restricted to 32kB due to nagios-core limit, which resulted in some long output being cut.

A workaround was implemented that stores output in files, which are exposed via link following

Full plugin output can be found at <link>

Same for WN tests. Additional links are now available in the JobSubmit output that point directly to job files such as jdl, stdout, log, env, for ARC gmlog can also be made available, etc.

Site alias	Local site etf
Hostname	arc-ce05.gridpp.rl.ac.uk
Service description	org.sam.CONDOR-JobSubmit-/cms/Role=lcgadmin
Service icons	
Service state	OK
Output of check plugin	OK - Job successfully completed
Long output of check plugin (multiline)	### Full plugin output can be found at /etf-raw/cms.Role.lcgadmin/scondor/arc-ce05.gridpp.rl.ac.uk/etf.log ### Job details (please follow links below to get additional debug information) ### Job jdl stdout log etf_log env === ETF job log: Timeout limits configured were: IDLE -> 1380 minutes global -> 1410 minutes Current time: 2020-11-02 10:28:09 Job started: 2020-11-02 10:13:08 Job finished: 2020-11-02 10:28:08 Job tracking times (entered): COMPLETED -> 2020-11-02 10:28:08 === Job JDL:

Already deployed in production for CMS, available to others if needed.

Timeouts

```
alice26.spbu.ru
org.sam.CREAMCE-JobSubmit-/hcb/Role=production
[icon]
CRIT
CRITICAL - Job (0) has failed with status: Job submission failed or timed out
=== Job JDL:
JDL(['(OutputSandBox', ['cream.out', 'wnlogs.tgz']), ('Executable', 'etf_run.sh'), ('StdError', 'cream.out'), ('StdOutput', 'cream.out'), ('OutputSandBoxBaseDes', ['var/lib/gridprobes/hcb.Role.production/cream/alice26.spbu.ru/etf_run.sh', 'var/lib/gridprobes/hcb.Role.production/cream/alice26.spbu.ru/gridjob.tgz']), ('J', 'JDL')])
=== Job submission:
glite-ce-job-submit --atm-delegation --debug -r alice26.spbu.ru/cream-pbs-lhcb /var/lib/gridprobes/hcb.Role.production/cream/alice26.spbu.ru/gridjob.jdl
2020-11-03 12:03:15,932 DEBUG - Using certificate proxy file [/opt/omd/sites/etf/etc/nagios/globus/userproxy.pem--lhcb]
2020-11-03 12:03:15,948 DEBUG - VO from certificate=[lhcb]
2020-11-03 12:03:15,948 WARN - No configuration file suitable for loading. Using built-in configuration
2020-11-03 12:03:15,948 DEBUG - Logfile is [/tmp/glite_cream_cli_logs/glite-ce-job-submit_CREAM_1002_20201103-120315.log]
2020-11-03 12:03:15,949 DEBUG - Processing file [/var/lib/gridprobes/hcb.Role.production/cream/alice26.spbu.ru/etf_run.sh]...
2020-11-03 12:03:15,949 DEBUG - Processing file [/var/lib/gridprobes/hcb.Role.production/cream/alice26.spbu.ru/gridjob.tgz]...
2020-11-03 12:03:15,949 DEBUG - Inserting mangled InputSandbox in JDL: [{"var/lib/gridprobes/hcb.Role.production/cream/alice26.spbu.ru/etf_run.sh"},
2020-11-03 12:03:15,953 INFO - certUtil:generateUniqueID() - Generated DelegationID: [c4d01efb8138732dc90c02a453b2471bd53999a6]
2020-11-03 12:03:16,101 FATAL - Connection to service [https://alice26.spbu.ru:8443/ce-cream/services/gridsite-delegation] failed: FaultString=[SSL error]
[SSL authentication failed in tcp_connect(): check password, key file, and ca file.]

Job submission has failed or timed out after 580 seconds.
Job submission failed or timed out
```

Site alias	Local site etf
Hostname	ce507.cern.ch
Service description	org.sam.CONDOR-JobSubmit-/cms/Role=logadmin
Service icons	[icon]
Service state	WARN
Output of check plugin	WARNING - Job (185916) has timed out while in status: RUNNING
Long output of check plugin (multiline)	### Full plugin output can be found at /etf-raw/cms.Role.lcgadmin/scondor/ce507.cern.ch/etf.log ### Job details (please follow links below to get additional debug information) ### Job: jdl stdout log etf_log env === ETF job log: Timeout limits configured were: IDLE -> 1380 minutes global -> 1410 minutes Current time: 2020-11-02 11:13:53 Job started: 2020-11-01 11:28:52 Job finished: 2020-11-02 11:13:52 Job tracking times (entered): RUNNING -> 2020-11-01 11:43:51

Submission timeout - submissions via HTCondor pool will always succeed even if resource is down, but will result in job being HELD due to GridResourceDown event after ~ 60 minutes.

Direct submissions timeouts are also supported (currently set to 10 minutes).

Once submitted, configurable timeouts on any job state - right now idle time and total time is being tracked - submission will result in warning (ETF job log has details)

Resubmission

alcyone-cms.grid.helsinki.fi				
State	Service	Icons	Status detail	
OK	org.cms.WN-analysis-/cms/Role=lcgadmin		alg19: OK	2020
OK	org.cms.WN-basic-/cms/Role=lcgadmin		alg19: OK	2020
OK	org.cms.WN-cvmfs-/cms/Role=lcgadmin		alg19: OK, cvmfs vers 2.7.2 (probe 1.3-pre2)	2020
OK	org.cms.WN-env-/cms/Role=lcgadmin		alg19: OK	2020
OK	org.cms.WN-frontier-/cms/Role=lcgadmin		alg19: OK	2020
OK	org.cms.WN-isolation-/cms/Role=lcgadmin		alg19: OK	2020
OK	org.cms.WN-mc-/cms/Role=lcgadmin		alg19: OK	2020
OK	org.cms.WN-squid-/cms/Role=lcgadmin		alg19: OK	2020
OK	org.cms.WN-xrootd-access-/cms/Role=lcgadmin		alg19: OK	2020
OK	org.cms.WN-xrootd-fallback-/cms/Role=lcgadmin		alg19: OK	2020
PEND	org.sam.CONDOR-JobReset-/cms/Role=lcgadmin			
OK	org.sam.CONDOR-JobState-/cms/Role=lcgadmin		Reschedule check	2020
OK	org.sam.CONDOR-JobSubmit-/cms/Role=lcgadmin		BI Aggregations containing this Service	2020
			Parameters for this service	

Currently testing in QA

In order to **allow sites to reset job submission**, which can sometimes get stuck (usually during/after downtime) - new metric will be added called JobReset.

Upon re-scheduling it will delete the current job file and re-submit a new job (so far this was only possible with manual intervention).

There is a time limit between two subsequent JobReset calls to avoid DoS (configurable, defaults to 15 minutes).

OK	org.sam.CONDOR-JobReset-/cms/Role=lcgadmin		OK - New job was submitted, please check JobState metric for progress	7.15 s	7.15 s
OK	org.sam.CONDOR-JobState-/cms/Role=lcgadmin		OK - Job was successfully submitted (194257)	2020-10-07 12:57:21	7.15 s

Worker Node micro-Framework (WN- μ FM)

- **Micro-framework to execute tests on the worker nodes**
 - Replaces statically compiled nagios binary, written in python with minimal dependencies
 - Supports py 2.6, 2.7, 3.4+; statically compiled version also available
- **Aims to provide reliable mechanism to run tests across different platforms**
 - Run tests in parallel (configurable), can timeout/kill runaway tests
 - Initial support for alternate schedule for WN tests (some tests don't need to run every time)
- **Runs nagios standard compliant tests, but also generic scripts**
 - Supports performance metrics (can report numeric values alongside status)
- **Configured directly from ETF**
 - Using env/config passed from JESS (via env file) - now easy to pass variables directly from frontend to the WN to be used in tests (like sitename, paths, originating CE/queue, etc.)
- **Pluggable support for test publishing**
 - Directory queues, message queues, json, http upload, etc.
 - Can also be used to parse script output and generate metrics from it
- **WN- μ FM can also run as a standalone component**

Framework is executed on the worker node using script `etf-run.sh` - details of tests execution available as part of the job stdout.

WN tests are part of the ETF experiments code base and can be executed separately. Improvements can be submitted via merge requests (see reference).

```
Nov 02 12:43:38 INFO core[13683]: *****
Nov 02 12:43:38 INFO core[13683]: * ETF WN micro Framework v.0.1.21: 2020-11-02 10:43:38.888962 *
Nov 02 12:43:38 INFO core[13683]: * Tests directory: /home/cmsgrid/sessiondir/AbQMdm78ltxnRDGs8pg5lGmPABFKDmABFKDmOvJKDmABFKDmwfizEm/etf/probes *
Nov 02 12:43:38 INFO core[13683]: * Output directory: /tmp/sam.13530.9982/msg-outgoing *
Nov 02 12:43:38 INFO core[13683]: * Python version: 2.6.6 *
Nov 02 12:43:38 INFO core[13683]: * Environment: *
Nov 02 12:43:38 INFO core[13683]: *   ETFROOT = /home/cmsgrid/sessiondir/AbQMdm78ltxnRDGs8pg5lGmPABFKDmABFKDmOvJKDmABFKDmwfizEm/etf *
Nov 02 12:43:38 INFO core[13683]: *   ETF_LEGACY = *
Nov 02 12:43:38 INFO core[13683]: *     etf/probes/org.cms/testjob/tests/CE-cms-env *
Nov 02 12:43:38 INFO core[13683]: *     etf/probes/org.cms/testjob/tests/CE-cms-squid.sing *
Nov 02 12:43:38 INFO core[13683]: *     etf/probes/org.cms/testjob/tests/CE-cms-analysis.sing *
Nov 02 12:43:38 INFO core[13683]: *     etf/probes/org.cms/testjob/tests/CE-cms-xrootd-access.sing *
Nov 02 12:43:38 INFO core[13683]: *     etf/probes/org.cms/testjob/tests/CE-cms-singularity *
Nov 02 12:43:38 INFO core[13683]: *     etf/probes/org.cms/testjob/tests/CE-cms-mc.sing *
Nov 02 12:43:38 INFO core[13683]: *     etf/probes/org.cms/testjob/tests/CE-cms-xrootd-fallback.sing *
Nov 02 12:43:38 INFO core[13683]: *     etf/probes/org.cms/testjob/tests/CE-cms-basic.sing *
Nov 02 12:43:38 INFO core[13683]: *     etf/probes/org.cms/testjob/tests/CE-cms-frontier.sing *
Nov 02 12:43:38 INFO core[13683]: *   ETF_TESTS = *
Nov 02 12:43:38 INFO core[13683]: *     etf/probes/org.cms/testjob/tests/CE-cms-env *
Nov 02 12:43:38 INFO core[13683]: *     etf/probes/org.cms/testjob/tests/CE-cms-squid.sing *
Nov 02 12:43:38 INFO core[13683]: *     etf/probes/org.cms/testjob/tests/CE-cms-analysis.sing *
Nov 02 12:43:38 INFO core[13683]: *     etf/probes/org.cms/testjob/tests/WN-cvmfs *
Nov 02 12:43:38 INFO core[13683]: *     etf/probes/org.cms/testjob/tests/CE-cms-xrootd: COMPLETED *
Nov 02 12:43:38 INFO core[13683]: *     etf/probes/org.cms/testjob/tests/CE-cms-singul: Worker node results tarball at: /var/lib/gridprobes/cms.Role.lcgadmin/scondor/arc04.lcg.csccs.ch/out/wnologs.tgz *
Nov 02 12:43:38 INFO core[13683]: *     etf/probes/org.cms/testjob/tests/CE-cms-mc.sing: Submitting worker node result: arc04.lcg.csccs.ch, org.cms.WN-frontier-/cms/Role=lcgadmin, OK (nid01894: OK) *
Nov 02 12:43:38 INFO core[13683]: *     etf/probes/org.cms/testjob/tests/CE-cms-xrootd: Submitting worker node result: arc04.lcg.csccs.ch, org.cms.WN-xrootd-fallback-/cms/Role=lcgadmin, OK (nid01894: OK) *
Nov 02 12:43:38 INFO core[13683]: *     etf/probes/org.cms/testjob/tests/CE-cms-basic.: Submitting worker node result: arc04.lcg.csccs.ch, org.cms.WN-basic-/cms/Role=lcgadmin, OK (nid01894: OK) *
Nov 02 12:43:38 INFO core[13683]: *     etf/probes/org.cms/testjob/tests/CE-cms-fronti: Submitting worker node result: arc04.lcg.csccs.ch, org.cms.WN-isolation-/cms/Role=lcgadmin, OK (nid01894: OK) *
Nov 02 12:43:38 INFO core[13683]: * ***** Submitting worker node result: arc04.lcg.csccs.ch, org.cms.WN-xrootd-access-/cms/Role=lcgadmin, OK (nid01894: OK) *
Nov 02 12:43:38 INFO core[13683]: * ***** Submitting worker node result: arc04.lcg.csccs.ch, org.cms.WN-mc-/cms/Role=lcgadmin, OK (nid01894: OK) *
Nov 02 12:43:38 INFO core[13683]: * ***** Submitting worker node result: arc04.lcg.csccs.ch, org.cms.WN-cvmfs-/cms/Role=lcgadmin, OK (nid01894: OK, cvmfs vers 2.7.3 (probe 1.3-pr2)) *
Nov 02 12:43:38 INFO core[13683]: * ***** Submitting worker node result: arc04.lcg.csccs.ch, org.cms.WN-analysis-/cms/Role=lcgadmin, OK (nid01894: OK) *
Nov 02 12:43:38 INFO core[13683]: * ***** Submitting worker node result: arc04.lcg.csccs.ch, org.cms.WN-squid-/cms/Role=lcgadmin, OK (nid01894: OK) *
Nov 02 12:43:38 INFO core[13683]: * ***** Submitting worker node result: arc04.lcg.csccs.ch, org.cms.WN-env-/cms/Role=lcgadmin, OK (nid01894: OK) *
Nov 02 12:43:38 INFO core[13683]: * ***** All WN results successfully submitted. *****
```

Short WN-FM summary is reported as part of the JobState metric output



Timeouts and errors - WN tests

```

gridce3.pi.infn.it
org.cms.WN-mc/cms/Role=lcgadmin
UNKN
n2wn21.pi.infn.it: UNKNOWN

### Full plugin output can be found at /etc/raw/cms.Role.lcgadmin/scondor/gridce3.pi.infn.it/CE-cms-mc.sing
Will run CE-cms-mc with Singularity if available
System: Linux n2wn21.pi.infn.it 3.10.0-1062.9.1.el7.x86_64 #1 SMP Fri Dec 6 15:49:49 UTC 2019 x86_64 x86_64
Current working directory: /home/grid/cmssgm/home_cream_171057416/CREAM171057416
drwxr-xr-x 3 cmssgm cms 136 Nov 3 12:40 .
SAME_SENSOR_HOME = /home/grid/cmssgm/home_cream_171057416/CREAM171057416/etf/probes/org
drwxr-xr-x 6 cmssgm cms 101 Nov 3 12:40 /home/grid/cmssgm/home_cream_171057416/CREAM171057416
PATH = /bin:/usr/local/lsf/bin:/usr/local/lsf/etc/sbin:/usr/sbin:/usr/bin:/usr/local/sbin:/home/grid/cmssgm/bin
lrwxrwxrwx 1 995 992 10 Sep 19 2017 /cvmfs/cms.cern.ch/SITECONF/local -> T2_IT_Pisa
Unsetting SCRAM_ARCH...
INFO Checking for singularity...
INFO Singularity at '/usr/bin/singularity' appears to work (privileged mode)
INFO Singularity found at '/usr/bin/singularity' (privileged mode, using s_bin)
OSG_SINGULARITY_VERSION = 2.5.0-dist
OSG_SINGULARITY_PATH = /usr/bin/singularity
OSG_SINGULARITY_EXTRA_OPTS = --home /home/grid/cmssgm/home_cream_171057416/CREAM171057416
Timezone set to UTC, be aware
Tue Nov 3 11:40:44 UTC 2020
uid=10251(cmssgm) gid=1023(cms) groups=1023(cms)
SCRAM_ARCH: slc6_amd64_gcc700
arch: slc6_amd64_gcc700
LD_LIBRARY_PATH:
/cvmfs/cms.cern.ch/COMP/slc6_amd64_gcc493/external/python/2.7.6/lib:/cvmfs/cms.cern.ch/COMP/slc6_amd64_gcc493/external/python/2.7.6/bin:/cvmfs/oasis.opensciencegrid.org/mis/osg-wn-client/3.4/3.4.55/el6-x86_64/lib64:/cvmfs/oasis.opensciencegrid.org/mis/osg-wn-client/3.4/3.4.55/el6-x86_64/X509_CERT_DIR:/cvmfs/oasis.opensciencegrid.org/mis/osg-wn-client/3.4/3.4.55/el6-x86_64/etc/grid-security/wmagent found python2.7 at...
/cvmfs/cms.cern.ch/COMP/slc6_amd64_gcc493/external/python/2.7.6/bin/python2.7

4
UNKNOWN - test timed out after 550 seconds
    
```

Better at catching test errors and reporting them
 - output now contains both stdout and stderr
 from test up to the point when test failed.



State	Service	Icons	Status detail
CRIT	org.atlas.WN-cvmfs-atlas/Role=lcgadmin		atlas-wn-210.roma1.infn.it: /usr/bin/stat: cannot read file system information for /cvmfs/atlas.cern.ch/repo/sw: No such file or directory
CRIT	org.atlas.WN-FrontierSquid-atlas/Role=lcgadmin		atlas-wn-210.roma1.infn.it: CRITICAL
CRIT	org.atlas.WN-swspace-atlas/Role=lcgadmin		atlas-wn-210.roma1.infn.it: CRITICAL
OK	org.sam.CONDOR-JobState-atlas/Role=lcgadmin		OK - Job was successfully submitted (1978869)
OK	org.sam.CONDOR-JobSubmit-atlas/Role=lcgadmin		OK - Job successfully completed


On WN total time and per test timeouts can be configured - if timeout occurs framework will attempt to retrieve output up to the point of the timeout. Test will be set to unknown.

Challenges and Plans

- **CREAM-CE retirement**
 - Proposing to stop testing in January - currently blocking update of the HTCondor pool
- **SRM storage metrics**
 - Legacy SRM tests retirement (moving to new protocols/gfal2 and python-nap needed)
- **Notifications**
 - Missing self-subscription functionality - would prefer a 3rd party component
- **Authentication**
 - Will need to re-implement to update to Checkmk 1.6/2.0 (we're now on 1.5) - looking for standalone container that would support new auth/authz (tokens/openid) but also x509
- **Resource landscape becoming complex (HPC, K8s, Clouds, etc.)**
 - Experiments workload management systems diverging in the way they use resources
 - In order to keep a common testing platform some form of integration with experiments WMS will be needed in the future
- **Short-term developments will focus on Checkmk 1.6/2.0 and K8s**
 - This will also include MW and C8 updates, moving to python3, etc.
 - All will be made available as part of continuous release process

Summary

- ETF is a container-based application combining open source software with a set of frameworks and APIs to provide flexible testing suite
- Easy to extend, re-locate and support new experiments and technologies
- Currently deployed at CERN for five experiments
 - Supporting IPv4-only and IPv6-only monitoring
 - Experiments contacts have access in case they need to debug and/or follow up on issues
 - Central instance provides a site-level view (one place to see results from all experiments)
- New job submission and WN framework offer a range of new possibilities
 - Easy to add new backends and more flexible ways for job submissions
 - New ways to run WN tests and collect/publish results
- Additional deployment at OSG for perfSONAR infrastructure monitoring
 - Strong interest from other communities to have this available as a generic tool
- MONIT reporting for IPv6 in development
- Feedback welcome via standard support channels or directly



Questions ?

Docs: <https://etf.cern.ch/docs/latest/> (to be updated soon)

Central instance: https://etf.cern.ch/etf/check_mk/

Instances (access requires IGTF/x509 cert loaded in the browser):

CMS production	CMS QA IPv6	CMS QA	Code: CMS gitlab
ATLAS production	ATLAS QA IPv6	ATLAS QA	Code: ATLAS gitlab
LHCb production		LHCb QA	Code: LHCb gitlab
ALICE production		ALICE QA	Code: ALICE gitlab
pS production		pS QA	Code: pS gitlab
		DUNE QA	Code: DUNE gitlab

ETF framework

[ETF core containers](#) [ETF Job Submission \(Jess\)](#)

[ETF nagios plugins lib. NAP](#) [ETF worker node micro framework](#)

ETF support channels: GGUS: Monitoring or etf-support@cern.ch (SNOW)

Backup slides

Architecture

ETF Core Framework

- Frontend API, configuration, scheduling, alerts

Plugins (probes/tests)

- Range of available plugins to test broad range of services
- Contributed by experiments, PTs, TFs and open source projects (Checkmk), etc.
- Python library to help write plugins (**python-nap**)

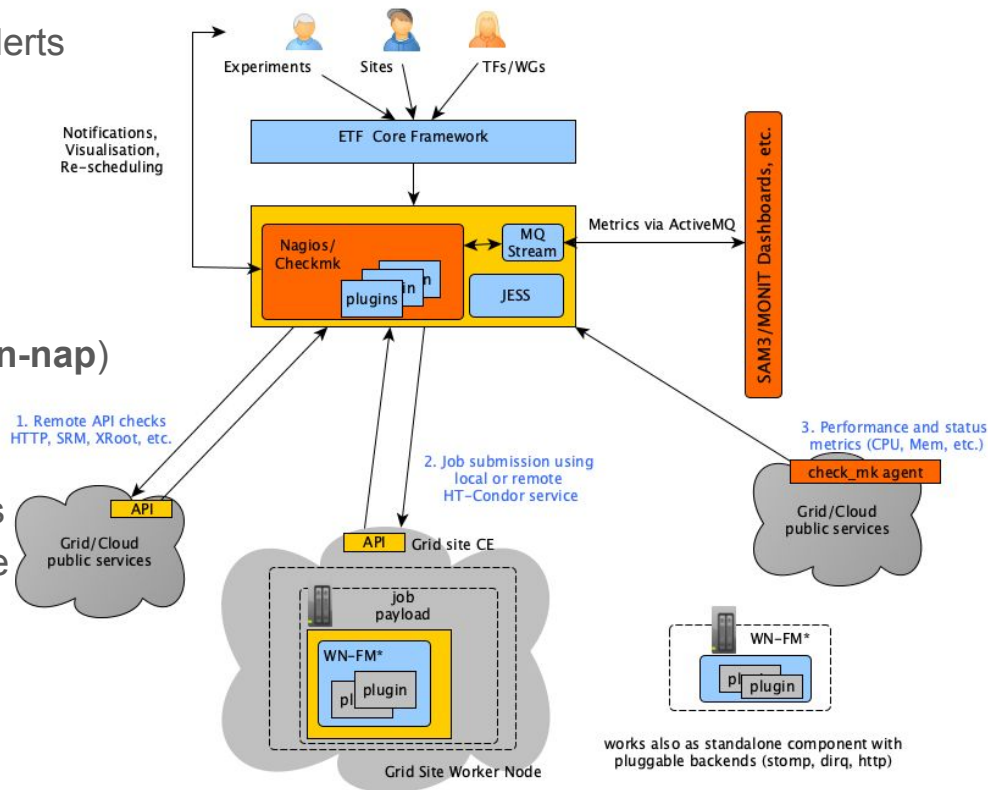
MQ Stream for publishing results

Job Submission Framework (JESS)

- Framework to write job submission plugins (submit/manage jobs, retrieve worker node results, etc.)

Worker Node Framework (WN-FM)

- Micro-scheduler to run tests on the WNs (configure and execute WN tests, collect results)




Deployment and Operations

Experiment instances @CERN (IPv4-only/IPv6-only in QA, IPv4-only in PROD)
perfSONAR infrastructure monitoring @OSG

ETF now runs in containers and uses gitlab CI with Auto DevOps

- Each experiment has its own container/image and gitlab repository
 - Full control over packages and versions to be deployed
- ETF can be deployed in the experiment-specific environment if needed
- Faster development cycle - changes propagated to QA upon each commit
 - Each commit triggers container rebuild and deployment to QA, one-click deploy to prod
- Simplified deployment - auto-deployed directly from gitlab
 - Easy to rollback



Core

ETF core currently running Checkmk 1.5

- Latest version is 1.6 (and 2.0alpha) - integration effort ongoing

ETF frontend - rule-based configuration system ([ncgx](#))

- **Pluggable** - python-based frontend/API for processing experiments topologies (what), tests (how) and schedule (when) ([example](#))
 - Template language to define how/when tests are executed
 - Tests can be configured to have different params per service/host/site, etc.
- Configuration now part of the experiments codebase (gitlab repository)
 - Experiments have full control over each aspect of the configuration ([including topology](#))
 - Gitlab repository - single place for code (images, tests) and configuration

Topology - hosts/service types currently taken from VO feeds

- XML-feed produced by the experiments ([docs](#), [intro](#))
 - Currently main source for hosts, queues, storage paths, etc.
 - While there are no limitations in ETF, for reporting/aggregation it's important to align topology with the aggregation layer (SAM3/MONIT)